

**This Page Is Inserted by IFW Operations
and is not a part of the Official Record**

BEST AVAILABLE IMAGES

**Defective images within this document are accurate representation of
The original documents submitted by the applicant.**

Defects in the images may include (but are not limited to):

- **BLACK BORDERS**
- **TEXT CUT OFF AT TOP, BOTTOM OR SIDES**
- **FADED TEXT**
- **ILLEGIBLE TEXT**
- **SKEWED/SLANTED IMAGES**
- **COLORS PHOTOS**
- **BLACK OR VERY BLACK AND WHITE DARK PHOTOS**
- **GRAY SCALE DOCUMENTS**

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

THIS PAGE BLANK (USPTO)

- 1 -

TITLE: METHOD AND SYSTEM FOR REMOTELY MONITORING TIME**VARIANT DATA****FIELD OF INVENTION**

The present invention relates to the field of remote data monitoring.

- 5 The invention has been developed primarily for use in remote monitoring and control of industrial plant, and will be described hereinafter with reference to this application. However, it will be appreciated by those skilled in the art that the invention is not limited to this field of use.

BACKGROUND

- 10 Remote monitoring of data from centralised observation areas is frequently used in industrial and manufacturing situations where it is desirable to closely monitor and, ideally, control, various systems and systems elements within those situations. The type of data to be monitored varies, and can include pressures, flow-rates, temperatures, items processed per unit time, or even access and security associated with a controlled area.
- 15 The use of one or more central observation areas reduces the need for direct observation of individual system elements by human operators. Moreover, the collection of data from multiple sources within a particular system enables instantaneous detection of faults or error conditions, as well as allowing processing of data from multiple sources to obtain better information about system operation at a higher level.
- 20 Remote monitoring of data requires considerable infrastructure development. For example, in the case of a gas production plant, it is typically necessary to provide many or even all major valves with transducers that measure desired information such as temperature and gas flow rates. The transducers are in turn connected to controllers that

- 2 -

can filter or otherwise manipulate the raw transducer data into a form useable by a remote monitoring server. The manipulation can include, for example, amplification and analog to digital conversion of the data from the transducer.

Frequently, the controllers are also used to control the components about which the transducers are measuring data. For example, a controller can filter, amplify and analog to digital convert data provided from a flow rate sensor adjacent a valve, whilst controlling throttling of the valve in accordance with instructions from a remote server. Other components within the system to be monitored will similarly require appropriate transducers to provide data to a central observation area.

Typically, the data is supplied to one or more data servers in digital form, the controller output representing the transducer being sampled periodically by the server and the resultant data values being stored in memory. In some cases, a predetermined number of samples are buffered by the server to ensure some historical data is available at all times to an operator. The data can be stored as a "point" on the server, the point simply being a structured object, containing accessible data. Usually, a point is stored in random access memory ("RAM") or on a hard disk drive associated with the server.

Frequently, it is desirable to use a relatively large-scale centralised control room to monitor (and sometimes control) many different areas of plant or even many plants at once. Where the different areas or plants are remote from each other, it may be necessary to implement one or more remote servers to collect data from them all. A coordinating server can then be used to collect data from the remote servers by means of a Wide Area Network ("WAN") or direct data link, for example.

- 3 -

A major disadvantage with this system is the need for complete duplication of information on both the coordinating and remote servers. As each remote server is brought into the system, all of the information about each and every point on that server must be supplied to the coordinating server. This enables the coordinating server to
5 handle a request from an operator to monitor a particular point or points on the remote servers. The coordinating server knows the correct address and parameters of each and every point on the network and can therefore access them as required.

Unfortunately, to ensure correct access to a point requested by an operator, it is necessary for the coordinating server to maintain current information in relation to each
10 and every point in the system. This means that for every new point and for every point for which a parameter changes, it is necessary to update the coordinating server so that it is aware of the change to the system. In a large system involving many components or even many remote plants, this can result in a relatively large amount of duplication of information on both the remote and coordinating servers, requiring correspondingly
15 large inputs of time on behalf of system administrators and programmers. In addition, the communication links between servers require high bandwidths in order to handle the increased traffic flow.

With a traditional system any node requiring data refers directly to the node that owns that data. The result is that nodes typically interrogate the owning node
20 independently, even if many nodes require the same data. When this scenario is extended to a connected sister system, or systems via a gateway or via other means, the situation just gets worse because there are simply more requestors of data.

- 4 -

An additional disadvantage of these prior art arrangements is that they require significant manual configuration of static data connections. Each time it is discovered that a system needs new data from another system, additional connections must be manually configured, resulting in increased engineering and maintenance. This also places a resource burden on each system, which has to mirror the full complement of all data required from other systems. While creating a bad situation with just two systems, connecting three or more systems becomes very unmanageable.

It is an object of the present invention to overcome or at least substantially ameliorate one or more of the disadvantages of the prior art.

10 It will be appreciated by those skilled in the art that the embodiments of the invention described in the detailed description provide other advantages over prior monitoring and control systems.

SUMMARY

According to a first aspect of the present invention there is disclosed a system for
15 remotely monitoring time variant data including:

a local server including a local server cache, a local controller, and a local operator station; and

at least one remote server, each remote server including a remote server cache and a remote controller; wherein:

20 the local server cache stores a plurality of data units;
the local controller provides updated data units to the local server cache;
each remote server cache stores a plurality of data units; and

- 5 -

each remote controller provides updated data units to its corresponding remote server cache.

According to a second aspect of the present invention there is disclosed a method of remotely monitoring time variant data using a system as defined in claim 1 wherein, when the local operator station wishes to subscribe to a selected data unit, the system performs the steps of:

(a) sending a selected data unit request from the local operator station to the local server; and

(b) determining whether the selected data unit is stored in the local server cache.

Preferably, the selected data unit is stored in the local server cache, and, having determined same, the system performs the step of:

(c) sending the selected data unit from the local server cache to the local operator station.

Alternatively, the selected data unit is not stored in the local server cache, and, having determined same, the system performs the step of:

(c) determining whether the selected data unit is provided by the local controller.

Preferably, the selected data unit is provided by the local controller, and, having determined same, the system performs the steps of:

(d) providing the selected data unit from the local controller to the local server cache; and

(e) sending the selected data unit from the local server cache to the local operator station.

- 6 -

Alternatively, the selected data unit is not provided by the local controller, and, having determined same, the system performs the step of:

(d) determining which remote controller the selected data unit is provided by.

Preferably, the selected data unit has a corresponding selected data unit identifier;

5 wherein the selected data unit is provided by a found remote controller;

and wherein the found remote controller is associated with

a found remote server, the server including:

a found remote server cache; and

a found remote server subscription list which stores the data unit identifiers

10 of the data units to which any server has subscribed and which are stored in the found remote server cache;

and, having determined that the selected data unit is provided by the found remote controller, the system performs the steps of:

(e) sending the selected data unit request from the local server to the found

15 remote server; and

(f) determining whether the selected data unit is stored in the found remote

server cache.

Preferably, the selected data unit is stored in the found remote server cache, and, having determined same, the system performs the steps of:

20 (g) sending the selected data unit from the found remote server cache to the local server cache; and

(h) sending the selected data unit from the local server cache to the local operator station.

- 7 -

Alternatively, the selected data unit is not stored in the found remote server cache, and, having determined same, the system performs the steps of:

(g) adding the selected data unit identifier to the found remote server subscription list;

5 (h) providing the selected data unit from the found remote controller to the found remote server cache;

(i) sending the selected data unit from the found remote server cache to the local server cache; and

(j) sending the selected data unit from the local server cache to the local operator
10 station.

According to a third aspect of the present invention there is disclosed a method of remotely monitoring time variant data by means of a subscribing server and one or more publishing servers interconnected by communications network hardware, the method including the steps of:

15 (a) establishing data points associated with the time variant data, each data point being disposed on one of the publishing servers and having a relatively unique identifier associated with it;

(b) providing each of the publishing servers on which a data point has been established with an address of the subscribing server;

20 (c) providing the subscribing server with the unique identifiers of the data points and addresses of the publishing servers; and

(d) selecting a first data point to be subscribed to by the subscribing server;

- 8 -

(e) searching the publishing servers for the unique identifier associated with the first data point; and

(f) upon locating the first data point, using the subscribing server to remotely monitor the time variant data associated therewith via the communications network hardware.

Preferably, upon locating the first data point, address data linking the first data point with the publishing server associated with it is recorded by the subscribing server, such that, upon subsequent selection of the first data point for remote monitoring, the first publishing server is the first of the publishing servers searched in step (e).

10 Preferably, as data points are located on publishing servers by the subscribing server, the subscribing server records at least the address of the publishing server on which that data point is disposed. Preferably, as data points are located, the corresponding publishing servers provide the subscribing server with other parameters associated with the respective data points.

15 In a second aspect, the present invention provide a system for remotely monitoring time variant data, the system including:

a plurality of publishing servers, upon each of which at least one data point associated with time variant data is established, each data point having a relatively unique identifier associated with it;

20 a subscribing server storing the unique identifiers of the data points and addresses of the publishing servers, each of the publishing servers including an address of the subscribing server;

- 9 -

a communications network interconnecting the subscribing and publishing servers;
and

data selection means enabling a user or other system to select a first data point to
be subscribed to by the subscribing server;

5 the system being configured such that, upon selection of a first data point via the
data selection means, the subscribing server searches the publishing servers to locate the
unique identifier associated with the first data point, and, upon locating the first data
point, the subscribing server remotely monitors the time variant data associated
therewith via the communications network.

10 Other exemplary aspects of the invention are set out in the numbered paragraphs at
the end of this specification.

BRIEF DESCRIPTION OF DRAWINGS

Preferred embodiments of the invention will now be described, by way of example
only, with reference to the accompanying drawings, in which:

15 Figure 1 is a schematic diagram of a system for remote monitoring time variant
data, in accordance with the invention;

Figure 2 is a detailed schematic diagram of one of the subscribing server shown in
Figure 1;

Figure 3 is a detailed schematic diagram of one of the publishing servers shown in
20 Figure 1;

Figure 4 is a schematic view of an exemplary gas processing plant to be monitored
and controlled by the system of the invention;

- 10 -

Figures 5 to 8 show examples of data points as supplied from publishing servers to a subscribing server in relation to the plant shown in Figure 4;

Figure 9 is a schematic diagram illustrating the basic operation of the present invention's caching strategy; and

Figure 10 is a schematic diagram illustrating how each server cache buffers controllers and other servers from unnecessary subscription transactions.

DETAILED DESCRIPTION

Referring to the drawings, there is shown a system 100 for remotely monitoring time variant data by means of a subscribing server 102 and a plurality of publishing servers 104. The subscribing server 102 and publishing servers 104 are interconnected by communications network hardware in the form of a Wide Area Network (WAN) 106.

It will be appreciated that the various servers can take any suitable form, in terms of both hardware and operating system specification. For example, one or more of the servers can take the form of an IBM compatible PC, Sparc Station, Apple Macintosh, or any other hardware platform. Similarly, the various platform or platforms utilised by the servers can run any suitable operating system, such as Windows 95, Windows NT, OS/2, Unix or the like.

Turning to Figure 2 the subscribing server 102 is shown in greater detail. It will be appreciated that many of the components required to operate a computer server are omitted for the sake of clarity, including elements such as the central processing unit ("CPU"), hard and floppy disk drives, CD-ROM and DVD drives, other storage media, output interfaces such as video displays and printers, and input interfaces such as keyboards and mouse or trackball controllers.

- 11 -

A memory area 200 of the subscribing server 102 stores a number of pieces of information allowing it to communicate effectively with the publishing servers 104 in accordance with the preferred embodiment. This information includes publishing server addresses 202 associated with the publishing servers 104 of interest. The publishing server addresses 202 can be supplied to the subscribing server 102 manually by means of an operator, or can alternatively be forwarded via the WAN by respective publishing servers 104 as they enter the system 100.

The memory area 200 also includes a database of data point identifiers 204 (data points and their identifiers are described below). Again, the data point identifiers 204 can be supplied manually directly to the subscribing server or can be provided across the WAN by respective publishing servers 104 as they enter the system 100. The subscribing server also includes one or more control terminals associated with it that can both extract data received by the subscribing server and send instructions via the subscribing server to the publishing server.

Turning to Figure 3, each publishing server 104 also includes a memory area 300, along with a processor and ancillary devices 302, which will typically be similar to those associated with the subscribing server 102. The memory area 300 initially includes one or more data points 304, each of which has at least one value associated with it. Data points can be considered objects in memory and can take the form of, for example, analog points (being digital representations of analog data), accumulator points and status points, or any known point structure. Each data point has time variant data associated with it. The data can be as simple as an alarm status or an instantaneous value associated with a remote transducer, as discussed below. In more complicated

- 12 -

embodiments, a point can represent an output of a function having any number of desired inputs. The inputs can include data from transducers or databases, or even values extracted from other data points.

In the preferred embodiment, at least some of the data points 304 relate to real world information supplied from corresponding transducers 306 which are linked to the publishing server 104 by means of appropriate sampling and filtering hardware in the form of controllers 308 that convert an output signal from a transducer into a value storable at a respective data point 304 located on the publishing server 104.

In use, when the system 100 is first established or initialised, in the preferred form there is no information linking particular data point identifiers 204, and thereby data points, with particular publishing servers 104. In the preferred embodiment, an operator of the subscribing server (or of a terminal associated therewith) is presented with a text field into which a data point identifier can be entered. A SEARCH/ RETRIEVE function key is pressed ("F12" in the preferred embodiment), and the subscribing server 102 then contacts each of the publishing servers in turn to find the location of that data point identifier. It will be appreciated that the precise method by which the desired data point identifier is selected on the subscribing server or associated terminal is not critical to the invention. For example, a directory or tree structure could be provided, hierarchically grouping different types of data point identifier together for selection by keystroke. In this and other embodiments, the required data point can be selected by use of a mouse or other pointing device.

Once the correct publishing server has been ascertained, a link is established with that publishing server to access the data point associated with the selected data point

- 13 -

identifier. In the preferred embodiment, once the correct publishing server has been identified, the subscribing server also requests information from the publishing server as to how the information from the data point of the unique data point identifier should be displayed to the operator of the subscribing server. For example, the publishing server
5 can indicate a particular page or template stored on the subscribing server suitable for displaying the information in the data point on the publishing server. Alternatively, the publishing server can supply the subscribing server with a suitable page or template, preferably using a generic or proprietary page description language.

Preferably, once the page or display template has been selected or received by the
10 subscribing server, the subscribing server uses field descriptors or other information in the page or template to ascertain what additional information needs to be requested from the publishing server to provide the operator with appropriate details of the data point. In yet other embodiments, all the information from the data point, including the live information, is supplied automatically to the subscribing server by the publishing server
15 upon which the desired data point has been found.

Once the publishing server upon which the data point of interest is located has been found, the subscribing server stores the mapping between the two within a cache 206 in the memory area 200. In this way, the next time that data point is requested by an operator of the subscribing server, the search for the data point can commence with the
20 mapped publishing server address in the cache 206. The cache can also be used to store current live data from the mapped publishing server, as well as the parameters and any metadata (both discussed below) associated with the data point. The cache can therefore provide a default model of the data point that can be updated the next time the data point

- 14 -

is accessed via the subscribing server. In addition, the mapped publishing server address in the cache 206 is deleted when the subscribing server detects that the data point has been deleted from the publishing server.

The publishing server can also store historical data in relation to its data points.

5 This historical data is then supplied to the subscribing data upon request. For example, upon receiving an alarm notification, the subscribing server or its operator may wish to ascertain conditions leading up to the event that caused the alarm to be generated. By requesting historical data in relation to the relevant data point or an element thereof, these historical conditions can be reconstructed.

10 It will be appreciated that one of the advantages provided by the invention, at least in preferred forms, is a reduction in data flow across the interconnecting network. This is because information is only sent between servers upon request, so information that is rarely accessed is only updated to the subscribing server when required.

Data points 304 can also define an alarm status associated with either a binary
15 output of a transducer, or a result of a function applied to a transducer output or data generated therefrom. For example, a reed switch installed to detect entry through a doorway provides a binary ON or OFF output that can be sent to a publishing server by an associated controller for storage at a data point for subsequent interrogation.

Alternatively, in the case of, for example, a transducer detecting a flow rate through a
20 pipe, a particular data point can raise an alarm flag within a predetermined data point if the flow rate exceeds a predetermined safe or desirable range of values.

Another aspect of the preferred embodiment is the ability to configure the system
100 to "push" data or alarm notifications from any of the publishing servers to the

- 15 -

subscribing server if predetermined conditions arise. For example, the subscribing server can instruct a particular publishing server to push data from a data point to the subscribing server in the event that a particular alarm flag is raised in response to predetermined conditions. The publishing server can also be instructed to push other
5 parameters associated with the data point being pushed.

A more specific embodiment of the invention will now be described in relation to Figure 4. A gas processing plant 400 includes two raw materials storage tanks 402 and 404 which supply raw materials to a mixing and heating tank 406 via respective gas control valves 401 and 403. Gas-powered heater 408 is disposed in the base of the
10 mixing and heating tank 406, and an agitator 410 in the form of a multi-bladed fan rotates within a central area of the mixing and heating tank.

As the raw gases are mixed, heated and processed, they rise to an upper area 405 of the mixing and heating tank 406, where they are extracted by pipe 412. Extraction via the pipe 412 is assisted by a pump 414 and controlled by a valve 415. The processed gas
15 is then moved to an interim-storage-tank 416 for later extraction.

Each of the valves and tanks, and the heater and the pump, has one or more transducers (not shown, for clarity) associated with it. The respective transducers are designed to provide specific real-time data in relation to a preselected aspect of its associated component's operation. For example, the valves in the present case have
20 transducers for flow rate, but in other embodiments measure parameters such as temperature, pressure, volume or any other desired characteristic of the gases flowing through them. Similarly, a transducer associated with the heater will usually measure temperature, whilst the pump's transducers measure the rotational rate of the pump and

- 16 -

the rate at which gas is being pumped. The particular transducers used in particular circumstances are well known within the art, and so are not discussed within this specification.

Each transducer will usually have a controller (again, not shown, for clarity) associated with it for amplifying and filtering the raw transducer output signal as required, as well as an analog to digital converter for placing the output into a form more suitable for storage in the memory of an associated publishing server. The controller then provides the processed signal from the transducer to a local publishing server.

Non time-variant data associated with the transducer, or the component it provides data on, can either be provided to the publishing server manually by an engineer or automatically by the controller at initialisation or some predetermined interval. Time variant data, usually the processed transducer output from the controller, is supplied at predetermined intervals. The intervals can be selected depending upon the likely rates of change of the transducer output, the importance of the data and the likely rates of change of the component property being measured by the transducer.

The data from the controller is stored on the server in an appropriate data point structure for retrieval by a subscribing server.

The controllers associated with the various transducers can also, in many cases, implement control instructions from the publishing server, or another instruction source. For example, if the mixed and heated gases in the mixing and heating tank are not being extracted quickly enough, the publishing server (or another instruction source) can instruct the controller associated with the pump 414 to increase the pump's speed.

- 17 -

The publishing server, and others like it, is connected to a communications network, which in turn is connected to a subscribing server. Operators of the subscribing server can access the information in the respective data points, as described above. For example, an operator may be interested in finding out information on the first raw materials storage tank 402, with the unique name T406. The unique name T406 is entered by the operator into a text field, and function key F12 (in this case) is pressed to initiate the process of searching the publishing servers on the network to find where the data point named T406 is stored. Once located, the publishing server instructs the subscribing server to use a specific template to display the data and then supplies the data point information. The specific template may be stored on the publishing server, or on the subscribing server, and is accessible via the subscribing server's display search path.

An example of the data point information at a particular point in time is shown in Figure 5. Values like CurrentVolume, CurrentTemperature and AlarmStatus are time variant, and are therefore supplied periodically to the subscribing server from the relevant publishing server. The operator at the subscribing server or an associated terminal can also observe the non time-variant data as required. In a preferred form, the rate at which the time-variant data is updated by the publishing server can be set by an operator of the subscribing server.

In a preferred embodiment, the operator at the subscribing server or an associated terminal can also control aspects of the gas processing plant using different fields in a data point being monitored. Preferably, there is a link between at least some of the time-variant elements of a data point and components of the gas processing plant which have

- 18 -

an effect on that element. For example, in the case of the mixing and heating tank's data point, one element is CurrentTemperature. Upon selecting this from the subscribing server, the subscribing server ascertains data points of other components that have an effect on the CurrentTemperature value of the mixing and heating tank data point. In this case, the heater is the main component of the gas processing plant that can affect the value of CurrentTemperature in the mixing and heating tank's data point. Accordingly, the heater's data point is searched for and located amongst the publishing servers, and then linked to the subscribing server to enable an operator to access the heater's controller. This linking of components that have an affect on the values of elements within other component's data points greatly simplifies the process of managing potentially complex systems and plant.

As discussed above, data points can also define an alarm status associated with either a binary output of a transducer, or a result of a function applied to a transducer output or data generated therefrom. In the example shown in Figure 4, an alarm data point can be generated at the publishing server associated with identifier T406, as shown in Figure 7. In this case, the alarm has arisen due to the temperature of the gas stored in the tank exceeding a predetermined temperature limit. It will be noted from the data point shown in Figure 5 that a HighTempLimit has been set at 280 kelvin and a HiHighTempLimit has been set at 300 kelvin. Once the temperature within the tank reaches the HighTempLimit of 280 kelvin, the alarm notification shown in Figure 7 is generated. It will be noted from the alarm notification that the temperature of the gas in the tank that generated the alarm was 290 kelvin. The alarm also records information of interest, including the date and time at which the alarm was generated, the relevant point

- 19 -

name, whether the alarm was acknowledged by an operator and whether the value returned to normal after the alarm was notified. It will be appreciated that any combination of these and other suitable parameters can be used.

It is desirable that alarm notifications are "pushed" from the relevant publishing server when they occur. Further, it is usual for information at each data points to be processed by their associated publishing server to ascertain whether an alarm notification is necessary. The way in which an alarm is delivered to the subscribing server can be set by the subscribing server itself, or by an operator at the publishing server.

In a preferred embodiment, once the subscribing server locates the publishing server upon which a data point of interest is located, metadata associated with that data point is supplied to the subscribing server to enable it to understand data subsequently supplied to it by the publishing server. For example, as shown in Figure 8, data point T4(X) includes metadata setting out each of the elements making up the data point, as well as the parameter data types, current quality of the parameters and, for example, a list of available modes for data supply. The subscribing server can use this information to understand how the data supplied in relation to the data point of interest can or must be used to ensure the operator is provided with a correct understanding of the status and values of the various data point elements. This information also informs the subscribing server of the various modes of operation it can place the data point into, as well as the types of data it can request if not supplied automatically.

An important aspect of at least a preferred embodiment of the invention is area handling, whereby different data points are grouped into logical clusters. This can be done upon instruction by the subscribing server, and enables control of access to

- 20 -

different types of data. For example, it may be desirable to restrict access to certain types of data, or the ability to change that data, to certain operators. A particularly advantageous aspect of this particular embodiment is the ability to combine different data points from different publishing servers into logical groups that can be accessed together or to which access can be limited on a operator by operator basis. Also, responsibility for a particular area can be handed over by simply reassigning the area, rather than all of the individual data points it contains.

It will also be appreciated that more than one subscribing server can be used within a network. Indeed, a particular server can simultaneously act as a publishing server for another subscribing server as well as a subscribing server to other publishing servers.

This interrelationship is perhaps more readily understood with reference to Figures 9 and 10.

In anticipation of the ever-increasing demand for information by users, the inventor used a next-generation approach from the outset by inserting a server-based cache between data users and the data owners. The cache forms a dynamic clearinghouse, constantly tailoring its contents based on current information demands on the system. In particular, subscriptions by multiple operator stations for common information are consolidated by the cache into a single subscription to the publishing server.

The server, as opposed to the controllers and control network, handles the direct impact of additional hits for data. This allows scalability to be increased significantly through the ever-increasing power of the PC servers. The controllers then simply publish subscribed data that is changing directly to the cache. The result is minimum

- 21 -

load on the controllers and control network.

Figure 9 illustrates the basic operation of the present invention's caching strategy.

There we can see a server cache 901 storing cached data 902 which it receives from a connected controller 903. The server (not shown) is connected to four operator stations 904, 905, 906 and 907. In this case, operator stations 904, 906 and 907 have subscribed to the data labelled FIC101.PV. In this subscribed state, the controller 903 publishes source data 908 about point FIC101.PV to the server cache 901 whenever that source data changes. The server cache 901 updates its cached data 902 and publishes that updated data to the three subscribing operator stations 904, 906 and 907. In this way, we can see that each of the subscribing operator stations does not interrogate the controller individually. The server cache 901 supplies the information to the three operator stations 904, 906 and 907 without having to interrogate the controller 903. The controller only supplies its source data 908 to the server cache 901 when that data changes, removing the supply of unnecessary, unchanged information.

Because only changed data is sent to the operator stations currently requiring it, the supervisory network sees greatly decreased traffic loading. This feature enables operator stations to offer unparalleled performance over remote dial-up links in a true client-server fashion. Prior art systems require the use of products (such as Carbon Copy or PC Anywhere) that tie up an entire operator station back at the plant for each remote user.

With the present invention, you dial into or connect over the Internet or Intranet to the system of choice, and enter your username and password. Your standard security profile is invoked independent of the location or computer you are working from - achieving true thin-client functionality.

- 22 -

With prior art systems, as has already been discussed, any node requiring data refers directly to the node that owns that data. The result is that nodes typically interrogate the owning node independently, even if many nodes require the same data.

In strong contrast to the architecture of prior art systems, the present invention is
5 able to integrate multiple servers and systems in a seamless fashion. The present invention builds on the caching architecture for stand-alone server illustrated in Figure 9 to provide transparent intersystem operations and information exchange with maximum performance and minimum impact among systems. The present invention allows applications and operator stations to subscribe not only to local points and alarms, but
10 also to points and alarms on other remote servers, all within a single global namespace. Global historical and real-time trending is also supported.

In the case of remote server-based data, the server subscribes with the remote data-owning server, which in turn subscribes with its connected controllers, unless the required data is already being subscribed by another source and is already in the cache.

15 Figure 10 shows an example of this process and how each server cache buffers controllers and other servers from unnecessary subscription transactions. The scenario depicted by the diagram is described step by step below.

Assume all the caches are empty initially

1. An operator station 1001 on Server C (not shown) subscribes to the "FIC101.PV"
20 data, causing Server C to subscribe to the FIC101.PV data in the locally attached controller 1005. The FIC101.PV data is now published on change to the Server C's cache 1006 and is propagated out to the subscribing operator station 1001.
2. Two additional operator stations 1003 and 1004 on Server C subscribe to FIC101.PV

- 23 -

data. Because FIC101.PV is already in the Server C's cache 1006, the two new operator stations 1003 and 1004 are simply added to the existing subscription list (not shown), and FIC101.PV is now published to all three subscribing stations 1001, 1003 and 1004 on Server C. No additional controller 1005 transactions are required.

5 3. Two operator stations 1007 and 1009 on Server A (not shown) now subscribe to FIC101.PV data. Because FIC101.PV is not in Server A's cache 1011, the system determines that FIC101.PV actually resides in Server C and subscribes to Server C. Because FIC101.PV is already in Server C's cache 1006, FIC101.PV data is now published on change to the three Server C operator stations subscribing it 1001, 1003
10 and 1004 as well as to Server A's cache 1011, which then publishes it to the two subscribing Server A operator stations 1007 and 1009. Again, no additional controller 1005 transactions are required.

4. Two operator stations on Server B (not shown) subscribe to "TI101.PV" data. The system determines that TI101.PV actually resides on Server C and subscribes to
15 Server C. However, the TI101.PV data is not already in the Server C cache 1006, so Server C is prompted to subscribe to the attached controller 1005 for it.

5. When the operator station 1009 on Server A also subscribes to Server C for TI101.PV, it is already in the Server C cache and is simply published on change to all subscribing parties.

20 In this example, everything occurs automatically, dynamically, and transparently. On the operator station, a user can type in the name of a remote point (such as FIC101.PV or TI101.PV), request its point detail, and dynamically the data is subscribed (in some cases hundreds of parameters), and the detail display is called up and filled with

- 24 -

current data. When the user or users are no longer interested in the data, the system de-subscribes, and, assuming that no other users are subscribing to the data, it is also dropped from the owning server's cache and from the owning controller's subscription list. This sequence of events occurs dynamically and without engineering intervention.

- 5 Complete transparency of point location makes all points appear as if they are resident on the local server from an application or operator's point of view.

Compared with the statically configured prior art arrangements, the present invention's infrastructure is more akin to an inter-server nervous system. As stimulus for data is received, the entire system reacts, dynamically setting up the linkages to transmit the required information. When the system senses that the information is no longer required, the linkage is automatically broken down, resulting in the minimum possible loading on system resources at all times.

The two-tiered caching structure (i.e., local server, remote server) of the present invention minimises traffic on control networks and intersystem local- and wide-area networks, as well as the supervisory network/Ethernet connecting operator stations. In the preferred embodiment, all data in the architecture is published on change, and then only to subscribing parties. Each server cache eliminates duplicate subscriptions, resulting in an N-to-1 reduction in network traffic and system loading at each level in the architecture.

20 In one preferred embodiment, the server cache includes features and performs functions substantially as disclosed in United States Patent No. 5,988,847 which was granted to Honeywell Inc. (Minneapolis) on 23 November 1999. The disclosure of that patent is incorporated herein by reference for the purpose of understanding more fully

- 25 -

the operation of the server cache. However, the disclosure of that patent is not admitted to be part of the prior art in relation to the present application for the purposes of novelty or inventiveness considerations.

The present invention provides a novel server architecture and method of
5 communication that provides substantial advantages over prior art remote monitoring and control systems. In particular, the "on-demand" nature of communication between servers reduces the average band width required of the interconnecting network. For this reason, the invention is particularly amendable to optimisation for low band width/WAN environments. In the preferred form, the use caching of recently accessed data points
10 and their associated addresses further increases the speed of the system and reduces unnecessary searching. The ability to assign data points from different servers to a single logical area similarly increases the flexibility of the system as a whole. For all of these reasons, the present invention represents a commercially significant advance over prior art, remote monitoring and control systems.

- 26 -

CLAIMS:-

1. A system for remotely monitoring time variant data including:

a local server including a local server cache, a local controller, and a local operator station; and

5 at least one remote server, each remote server including a remote server cache and a remote controller; wherein:

said local server cache stores a plurality of data units;

said local controller provides updated data units to said local server cache;

each remote server cache stores a plurality of data units; and

10 each remote controller provides updated data units to its corresponding remote server cache.

2. A system as claimed in claim 1 wherein said local operator station uses or views said data units stored in said local server cache one of:

(a) whenever it requests said data;

15 (b) whenever said data units change;

(c) at predetermined intervals; or

(d) at operator station specified intervals.

3. A system as claimed in claim 1 wherein said local controller provides said updated data units to said local server cache one of:

20 (a) whenever said data units change;

(b) at predetermined intervals; or

(c) at operator station specified intervals.

- 27 -

4. A system as claimed in claim 1 wherein each remote controller provides said updated data units to its corresponding remote server cache one of:

- (a) whenever those data units change;
- (b) at predetermined intervals; or
- 5 (c) at operator station specified intervals.

5. A system as claimed in claim 1 wherein:

said local server further includes a local server subscription list which stores the data unit identifiers of the data units to which any other server has subscribed and which are stored in said local server cache.

10 6. A system as claimed in claim 5 wherein:

said local server cache publishes said local cache-stored data units to the server caches so indicated in said local server subscription list, one of:

- (a) whenever those data units change;
- (b) at predetermined intervals; or
- 15 (c) at operator station specified intervals.

7. A system as claimed in claim 1 wherein said remote server further includes a remote server subscription list which stores the data unit identifiers of the data units to which any other server has subscribed and which are stored in said remote server cache

8. A system as claimed in claim 7 wherein said remote server cache publishes said
20 remote cache-stored data units to the server caches so indicated in said remote server subscription list, one of:

- (a) whenever those data units change;
- (b) at predetermined intervals; or

- 28 -

(c) at operator station specified intervals.

9. A system as claimed in claim 1 wherein new servers may be readily added to said system.

10. A system as claimed in claim 1 wherein servers may be readily disconnected from
5 said system.

11. A system as claimed in claim 1 wherein each data unit has metadata associated with it.

12. A system as claimed in claim 11 wherein said metadata includes one or more of:

(a) data content information; and

10 (d) data display information.

13. A system as claimed in claim 12 wherein said data content information includes information relating to one or more of:

(a) the names of the parameters making up said data unit;

(b) the data types of the parameters making up said data unit; and

15 (c) the current quality of said data.

14. A system as claimed in claim 12 wherein said data display information indicates how said data unit should be displayed on an operator station and includes one or more of:

(a) a page template; or

20 (b) a display template.

15. A system as claimed in claim 14 wherein a selected data unit has corresponding selected data display information and wherein that information is stored on a found server which owns a found controller which supplies said selected data unit.

- 29 -

16. A system as claimed in claim 15 wherein said selected data display information is accessible via a display search path on the operator station's server.

17. A system as claimed in claim 14 wherein said display template includes field descriptors

5 18. A system as claimed in any one of claims 11 to 17 wherein, when a selected data unit is first subscribed to by said local operator station, said system forwards said selected data unit's metadata to said local operator station.

19. A system as claimed in claim 1 wherein said data units include one or more of:

(a) data points;

10 (b) notifications; and

(c) historical data.

20. A system as claimed in claim 19 wherein said notifications include one or more of:

(a) alarm notifications;

(b) messages; and

15 (c) alarm confirmations.

21. A system as claimed in claim 19 wherein said data points include one or more of:

(a) analog points;

(b) accumulator points;

(c) status points;

20 (d) control modules; and

(e) any other point structure

- 30 -

22. A system as claimed in claim 19 wherein said local controller is in communication with a local transducer and wherein said data points contain real world information provided by said transducer.

23. A system as claimed in claim 22 wherein said real world information includes time variant data.

24. A system as claimed in claim 23 wherein said time variant data includes one or more of:

- (a) an alarm status;
- (b) an instantaneous value associated with a transducer; or
- (c) an output of a function having a plurality of inputs.

25. A system as claimed in claim 24 wherein said inputs include one or more of:

- (a) data from said transducer;
- (b) data from other databases; and
- (c) data extracted from other data points.

26. A system as claimed in claim 24 wherein said alarm status has at least an active state and an inactive state, said active state being indicative of an undesirable value associated with a data point and said inactive state being indicative of a desirable value associated with a data point.

27. A system as claimed in claim 26 wherein said alarm status changes from said

inactive state to said active state as a result of one or more of:

- (a) an output of said transducer;
- (b) a function applied to an output of said transducer; or
- (c) data generated from said transducer.

- 31 -

28. A system as claimed in claim 19 wherein said system monitors said data points in order to ascertain whether an alarm notification is necessary.

29. A system as claimed in claim 28 wherein, when an alarm notification is necessary said system publishes said alarm notification to all servers which subscribe to the data point.

30. A system as claimed in claim 19 wherein said servers store historical data relating to at least one of said data points disposed on that server.

31. A system as claimed in claim 1 wherein a subset of said plurality of data units are grouped into logical clusters.

10 32. A system as claimed in claim 31 wherein said logical clusters include data units stored on a plurality of different servers.

33. A system as claimed in claim 31 wherein access to said logical clusters is restricted to particular human operators or classes of human operators.

34. A system as claimed in claim 1 wherein said local and remote servers are linked
15 by communications network hardware.

35. A system as claimed in claim 34 wherein said communications network hardware includes low bandwidth communication links.

36. A system as claimed in claim 35 wherein said low bandwidth communications links includes WAN or LAN communication links.

20 37. A system as claimed in claim 1 wherein said local server further includes a local server memory area.

38. A system as claimed in claim 37 wherein said local server memory area stores one or more of:

- 32 -

(a) data point addresses; and

(b) data point identifiers.

39. A system as claimed in claim 1 wherein said remote server further includes a remote server memory area.

5 40. A system as claimed in claim 39 wherein said remote server memory area stores one or more of:

(a) data point addresses; and

(b) data point identifiers.

10 41. A system as claimed in claim 1 wherein said data units are provided to said controllers by a plurality of time variant data sources.

42. A system as claimed in claim 41 wherein each time variant data source is monitored by a corresponding transducer.

15 43. A system as claimed in claim 41 wherein said controller includes sampling and filtering hardware that converts an output signal from said transducer into a value storable at its corresponding data point.

44. A system as claimed in claim 41 wherein said servers control said time variant data sources by sending controller instructions to said controller which corresponds to that time variant data source.

20 45. A system as claimed in claim 44 wherein at least one of said time variant data sources has an associated effecting component which effects the time variant data of said data source.

- 33 -

46. A system as claimed in claim 45 wherein said system controls said time variant data source achieved by sending said controller instruction to that data source's associated effecting component.

47. A system as claimed in claim 1 wherein when said system is first established or
5 initialised, said caches do not contain any data units.

48. A method of remotely monitoring time variant data using a system as claimed in any one of claims 1 to 47 wherein, when said local operator station wishes to subscribe to a selected data unit, said system performs the steps of:

(a) sending a selected data unit request from said local operator station to said
10 local server, and

(b) determining whether said selected data unit is stored in said local server cache.

49. A method as claimed in claim 48 wherein said selected data unit is stored in said local server cache, and, having determined same, said system performs the step of:

15 (c) sending said selected data unit from said local server cache to said local operator station.

50. A method as claimed in claim 48 wherein said selected data unit is not stored in said local server cache, and, having determined same, said system performs the step of:

(c) determining whether said selected data unit is provided by said local
20 controller.

51. A method as claimed in claim 50 wherein said selected data unit is provided by said local controller, and, having determined same, said system performs the steps of:

- 34 -

(d) providing said selected data unit from said local controller to said local server cache; and

(e) sending said selected data unit from said local server cache to said local operator station.

5 52. A method as claimed in claim 50 wherein said selected data unit is not provided by said local controller, and, having determined same, said system performs the step of:

(d) determining which remote server said selected data unit is provided by.

53. A method as claimed in claim 52 wherein said selected data unit has a corresponding selected data unit identifier; wherein said selected data unit is provided by

10 a found remote server; and wherein

said found remote server including:

a found remote server cache; ;

a found remote controller and

15 a found remote server subscription list which stores the data unit identifiers of the data units to which any server has subscribed and which are stored in said found remote server cache;

and, having determined that said selected data unit is provided by said found remote server, said system performs the steps of:

20 (e) sending said selected data unit request from said local server to said found remote server; and

(f) determining whether said selected data unit is stored in said found remote server cache.

- 35 -

54. A method as claimed in claim 53 wherein said selected data unit is stored in said found remote server cache, and, having determined same, said system performs the steps of:

(g) sending said selected data unit from said found remote server cache to said
5 local server cache; and

(h) sending said selected data unit from said local server cache to said local operator station.

55. A method as claimed in claim 53 wherein said selected data unit is not stored in said found remote server cache, and, having determined same, said system performs the
10 steps of:

(g) adding said selected data unit identifier to said found remote server subscription list;

(h) providing said selected data unit from said found remote controller to said found remote server cache;

15 (i) sending said selected data unit from said found remote server cache to said local server cache; and

(j) sending said selected data unit from said local server cache to said local operator station.

56. A method of remotely monitoring time variant data using a system as claimed in
20 any one of claims 1 to 47 wherein said system performs the steps of:

(a) scanning said local server cache to identify those data units which have not been used or viewed by said local operator station in a predetermined period; and

(b) removing those data units from said local server cache.

- 36 -

57. A method of remotely monitoring time variant data using a system as claimed in any one of claims 1 to 47 wherein said system performs the steps of:

(a) scanning each remote server cache to identify those data units which have not been used or viewed by any operator station in a predetermined period; and

5 (b) removing those data units from each remote server cache.

58. A method of remotely monitoring time variant data using a system as claimed in any one of claims 1 to 47 wherein, when a new server is added to said system, the address of said new server is manually provided to all of the existing servers in the system.

10 59. A method as claimed in claim 48 wherein said step of sending a selected data unit request from said local operator station to said local server is preceded by said local operator station receiving a data unit identifier from a human operator.

60. A method of remotely monitoring time variant data substantially as hereinbefore described with reference to the accompanying drawings and examples.

15 61. A system for remotely monitoring time variant data substantially as hereinbefore described with reference to the accompanying drawings and examples.

62. A method of remotely monitoring time variant data by means of a subscribing server and one or more publishing servers interconnected by communications network hardware, the method including the steps of:

20 (a) establishing data points associated with the time variant data, each data point being disposed on one of the publishing servers and having a relatively unique identifier associated with it;

- 37 -

(b) providing each of the publishing servers on which a data point has been established with an address of the subscribing server;

(c) providing the subscribing server with the unique identifiers of the data points and addresses of the publishing servers; and

5 (d) selecting a first data point to be subscribed to by the subscribing server;

(e) searching the publishing servers for the unique identifier associated with the first data point; and

(f) upon locating the first data point, using the subscribing server to remotely monitor the time variant data associated therewith via the communications
10 network hardware.

63. A method according to claim 62, wherein, upon locating the first data point, address data linking the first data point with the publishing server associated with it is recorded by the subscribing server, such that, upon subsequent selection of the first data point for remote monitoring, the first publishing server is the first of the
15 publishing servers searched in step (e).

64. A method according to claim 62 or 63, wherein step (c) includes substep of:

(ca) as each publishing server is established, providing that publishing server's address to the subscribing server via the communications network hardware.

65. A method according to any one of claims 62 to 64, wherein at least one of the
20 data points is an alarm point having at least an active state and an inactive state, the active state being indicative of an undesirable value associated with a data point and the inactive state being indicative of a desirable value associated with a data point.

66. A method according to claim 65, further including the step of:

- 38 -

(h) upon the alarm point entering the active state, pushing an alarm status indicator to the subscribing server via the communications network hardware.

67. A method according to claim 66, wherein step (h) includes a substep of:

(ha) pushing data associated with the alarm point to the subscribing server via
5 the communications network hardware.

68. A method according to claim 67, wherein the data associated with the alarm point includes the data for which an undesirable value has been generated.

69. A method according to any one of claims 62 to 68, wherein step (f) includes a substep of:

10 (fa) instructing the publishing server associated with the selected data point to provide properties associated with the selected data point to the subscribing server.

70. A method according to claim 69, wherein the properties referred to in substep (fa) include metadata associated with the data point.




FIG. 1

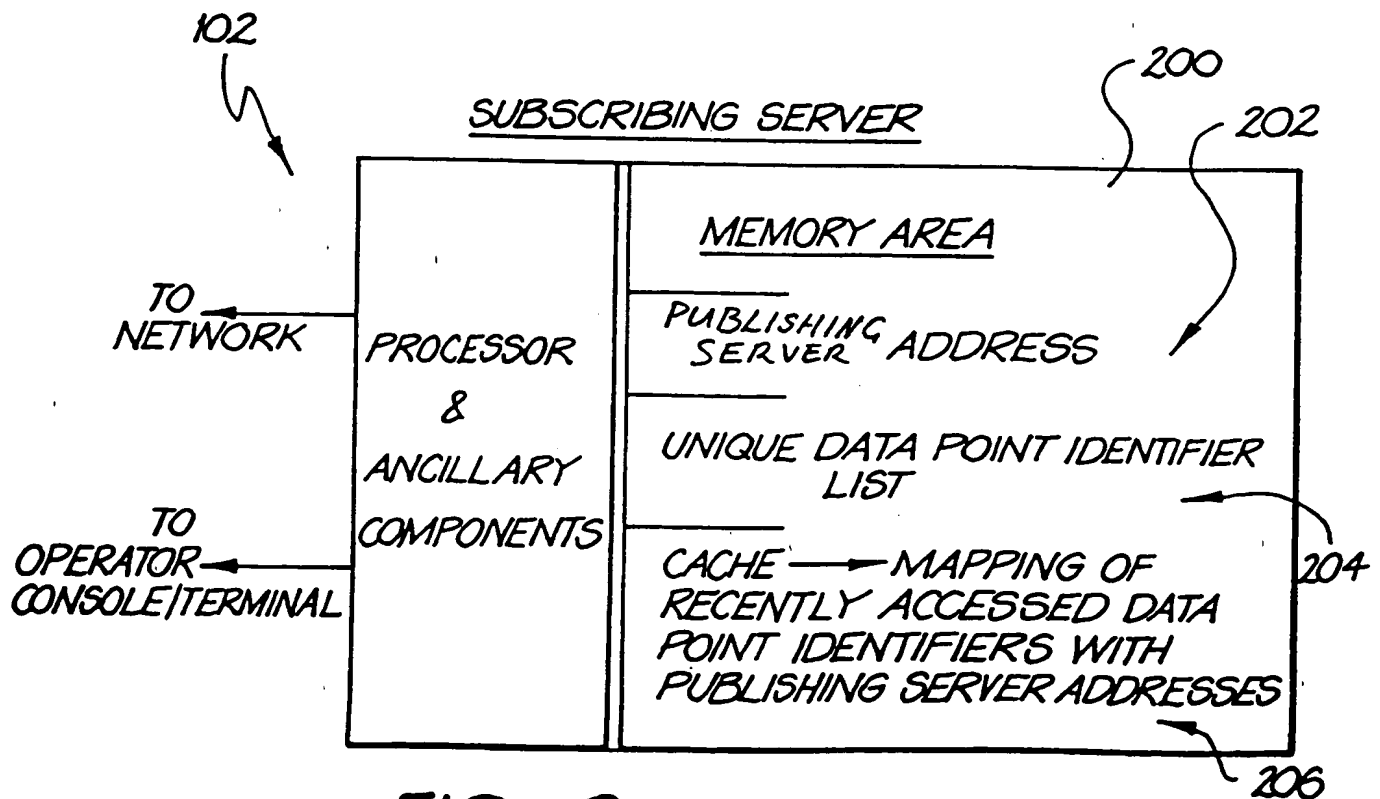



FIG. 2



3/7




FIG. 4

Figure 5 - Point T406

<i>Point Property</i>	<i>Value</i>
Name	T406
Description	Methane Storage Tank
Area	A1
CreateTime	10:20 1/3/99
CurrentVolume	31000.23 cubic meters
CurrentTemperature	290 Kelvin
AlarmStatus	HiTemp
HighTempLimit	280 Kelvin
HiHighTempLimit	300 Kelvin
DetailDisplayName	StorageTank.dsp
GroupDisplayName	StorageTankGroup.dsp

Figure 6 - Point V401

<i>Point Property</i>	<i>Value</i>
Name	V401
Description	Methane outlet control valve
Area	A1
CreateTime	10:20 1/3/99
PresentValue	22.23 cubic meters per hour
SetPoint	22.34 cubic meters per hour
DetailDisplayName	ButterFlyValve.dsp
GroupDisplayName	ButterFlyValveGroup.dsp
Mode	Automatic
Output	20%

Figure 7 - Alarm from T406

<i>Alarm Property</i>	<i>Value</i>
DateTime	10:03PM 9/6/99
PointName	T406
Priority	Urgent
AlarmType	HighTemperature
Value	290
Acknowledged	No
ReturnToNormal	No
Limit	280

Figure 8 - METADATA for T406

<i>Request</i>	<i>Value</i>
Get parameter names	Name, Description, Area, CreateTime, CurrentVolume, CurrentTemperature, AlarmStatus, HighTempLimit, HiHighTempLimit, DetailDisplayName, GroupDisplayName
Get parameter data types	String, String, String, Date, Number, Number, String, Number, Number, String, String
Get quality of parameters	Ok,ok,ok,ok,ok,ok,ok,ok,ok,ok,BAD, InAlarm
List enumerations of mode	Automatic, Manual, Cascade
Get range of Current Temperature	0 to 400
Get how much to raise Current Temperature	+10 increments
Get length of Name	4 character

6/7




FIG. 9

7/7

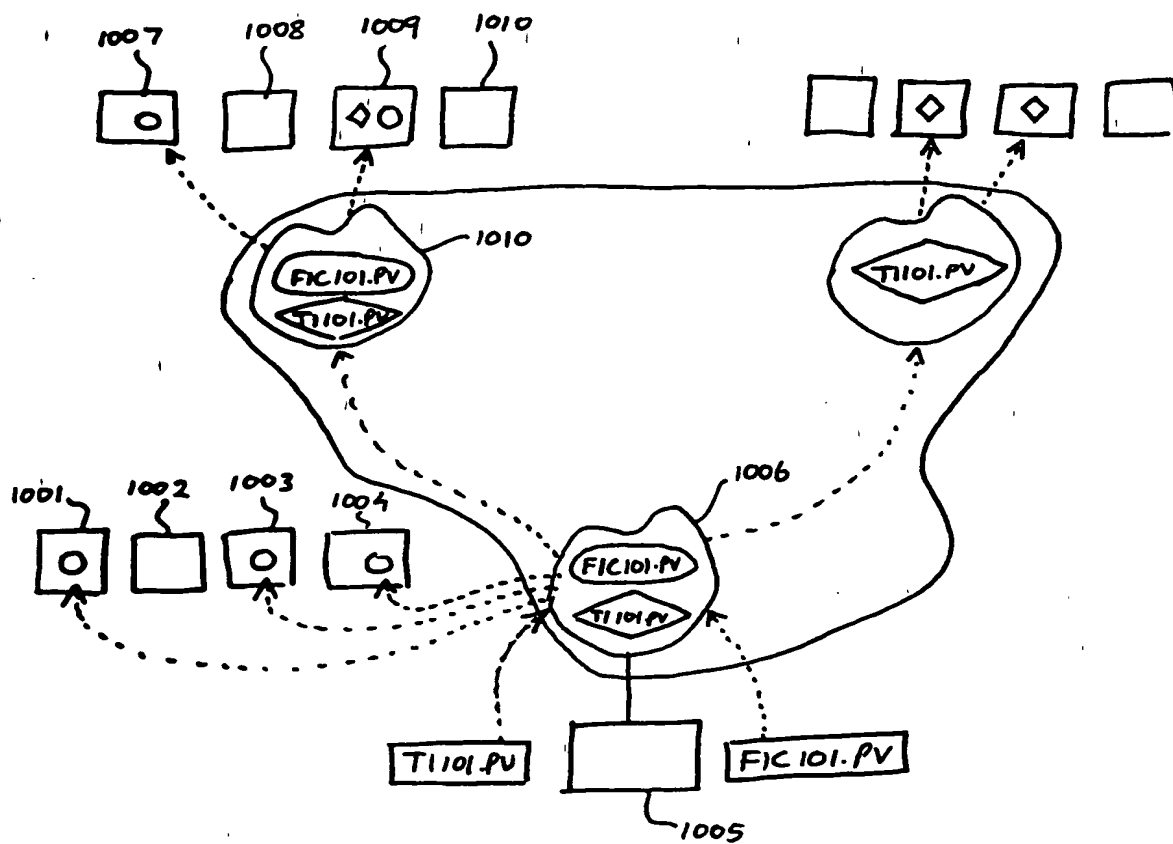


FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU00/00658

A. CLASSIFICATION OF SUBJECT MATTERInt. Cl. ⁷: G05B 15/00, 15/02, 19/048; G06F 15/16, 17/60

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC G05B 15/00, 15/02, 19/00, 19/048, G06F 15/16, 17/60, 19/00, F28F 27/00, F28F 27/02

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
AU: IPC AS ABOVE

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPAT: SERVER? OR CLIENT? OR HOST?; MONITOR? OR SUPERVIS+ OR CONTROL+ OR
SURVEILLANCE; COMPUTER?;Industr?; MANUFACTUR+; FACTORY OR FACTORIES;
PROCESS+; PLANT?; WAN OR WIDE AREA NETWORK?; TRAFFIC FLOW**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X, Y	US 5862391 A (SALAS et al.) 19 January 1999 Entire document, particularly columns 1 - 12, column 16 line 62 to column 17 line 40, claims and drawings	1 - 70
X, Y	US 5463735 A (PASCUCCHI et al.) 31 October 1995 Entire document	1 - 70

☒ Further documents are listed in the continuation of Box C ☒ See patent family annex

* Special categories of cited documents:

"A" Document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

9 August 2000

Date of mailing of the international search report

18 August 2000

Name and mailing address of the ISA/AU

AUSTRALIAN PATENT OFFICE
PO BOX 200, WODEN ACT 2606, AUSTRALIA
E-mail address: pct@ipaustalia.gov.au
Facsimile No. (02) 6285 3929

Authorized officer

CHARLES BERKO
Telephone No : (02) 6283

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU00/00658

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5862346 A (KLEY et al.) 19 January 1999	1, 62
Y	Entire document	2-61, 63-70
P, X	US 6088717 A (REED et al.) 11 July 2000	1, 62
P, Y	Entire document	2-61, 63-70
P, X	US 5987621 A (DUSO et al.) 16 November 1999	1, 62
P, Y	Entire document	2-61, 63-70
P, Y	US 5988847 A (MCLAUGHLIN et al.) 23 November 1999	1 - 70
Y	US 5774661 A (CHATTERJEE et al.) 30 June 1998	1, 62
A	Entire document	2-61, 63-70
Y	US 5862054 A (LI) 19 January 1999	1, 62
A	Entire document	2-61, 63-70
Y	WO 98/55934 A (ALPHA TECHNOLOGIES INC.) 10 December 1998	1, 62
A	Entire document	2-61, 63-70

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/AU00/00658

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report			Patent Family Member				
US	5862391	NONE					
US	5463735	AU	73304/91	CA	2075048	EP	513206
		JP	7182283	JP	8055051	WO	9111766
		US	5384697	US	5444851	US	5511188
		US	5522044	US	5550980	US	5598566
		US	5884072				
US	5862346	NONE					
US	6088717	NONE					
US	5987621	US	5892915				
US	5988847	AU	91058/98	EP	1004064	WO	9910787
US	5774661	NONE					
US	5862054	NONE					
WO	9855934	AU	80571/98	US	5961604		
END OF ANNEX							

THIS PAGE BLANK (USPTO)